



In re Application of:

Thomas W. Lantzatella, et al.

Serial No. 09/997,602

Filed: November 29, 2001

For: METHODS,
FUNCTIONAL DATA,
AND SYSTEMS TO
REPRESENT A STORAGE
ENVIRONMENT

§ Group Art Unit: 2168

§ Examiner: Alaubaidi, Haythim J.

§ Atty. Dkt. No.: 5760-16700

I hereby certify that this correspondence is being deposited with the U.S. Postal Service with sufficient postage as First Class Mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date indicated below:

B. Noël Kivlin

Name of Registered Representative

February 21, 2006
Date

Signature

Mail Stop Appeal Brief - Patents

Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Appeal filed December 19, 2005, Appellants present this Appeal Brief. Appellants respectfully request that the Board of Patent Appeals and Interferences consider this appeal.

02/28/2006 HAHMED1 00000095 501505 09997602

01 FC:1402 . 500.00 DA

BEST AVAILABLE COPY

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.

I. REAL PARTY IN INTEREST

The subject application is owned by VERITAS Operating Corporation, a corporation organized and existing under and by virtue of the laws of the State of Delaware, and now having its principal place of business at 350 Ellis Street, Mountain View, CA 94043.

II. RELATED APPEALS AND INTERFERENCES

No other appeals, interferences or judicial proceedings are known which would be related to, directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1-26 stand finally rejected. The rejection of claims 1-26 is being appealed. A copy of claims 1-26 is included in the Claims Appendix herein below.

IV. STATUS OF AMENDMENTS

No amendments to the claims have been submitted subsequent to the final rejection.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 is directed toward a method to map a storage environment data object. The method comprises receiving a reference to the data object in a first storage environment, wherein the data object resides in a second storage environment, and generating a first data structure from the reference representing one or more physical locations of the data object within the second storage environment. In addition, the method comprises associating a signature with the data object, wherein the signature is indicative of a state of the data object, and retaining the first data structure in the first

storage environment. The method further comprises querying the second storage environment for a change to the signature in preparation for a data access operation on the data object, and updating the first data structure if the signature has changed. The method also includes performing the data access operation using the first data structure to interface with one or more of the physical locations of the data object from the first storage environment. *See, e.g.,* Fig. 1, Fig. 2, Fig. 3 and Fig. 5; page 8, lines 15 – 30; page 9, lines 15 - 30; page 10, lines 16 – 30; and page 12, line 3 – page 13, line 6.

Independent claim 9 is directed to a method to represent a data storage object. The method includes identifying one or more storage locations for the data storage object housed within a first storage environment and assembling a hierarchical map representing a path to one or more of the storage locations. The method also includes associating a signature with the map indicative of a state of the data storage object and querying the signature for changes in preparation for a data access operation on the data storage object. The method further comprises updating the map if the signature has changed, and using the map in a second storage environment to access the data object. *See, e.g.,* Fig. 1, Fig. 3, Fig. 4 and Fig. 5; page 8, lines 15 – 30; page 9, lines 1 - 14; page 10, lines 16 – 30; and page 14, line 11 – page 15, line 14.

Independent claim 15 is directed towards a first computer readable medium having a data map and a signature representing a data object residing on a second computer readable medium, wherein the signature is indicative of a state of the data object. The map comprises a plurality of nodes, including a first node representing the data object, a file system node representing a file system on the second computer readable medium, a volume node representing a volume manager associated with the file system, one or more partition nodes managed by the volume manager, and one or more disk identifications representing one or more storage devices housing the data object. The map is updated when a change to the signature is detected. *See, e.g.,* Fig. 4 and Fig. 5; page 7, line 13 – page 8, line 7; page 10, lines 16 – 30; and page 15, line 19 – page 17 line 24.

Independent claim 21 is directed to a storage environment system. The storage environment system comprises a first and a second file system. One or more data objects residing in the second file system are capable of being referenced within the first file system. The system also includes a map and a signature. The map is generated to represent one or more physical locations for each of the one or more data objects and used by the first file system when one or more of the data objects are referenced. The signature indicates a state of the one or more data objects. The map is updated when changes are detected and associated with the signature. *See, e.g.*, Fig. 1 and Fig. 5, page 10, lines 16 – 30; and page 18, lines 1 – 22.

VI. GROUND S OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1 – 25 are rejected under 35 U.S.C. 103(a) as being unpatentable over Frey et al. (U.S. Patent Number 6,029,168, hereinafter ‘Frey’) in view of Mukherjee (U.S. Patent Number 6,466,978, hereinafter ‘Mukherjee’).

2. Claim 26 is rejected under 35 U.S.C. 103(a) as being unpatentable over Frey in view of Mukherjee, further in view of Banerjee et al., (U.S. Patent No. 6,795,830, hereinafter, ‘Banerjee’).

VII. ARGUMENT

First Ground of Rejection:

Claims 1 - 25 stand finally rejected under 35 U.S.C. § 103(a) as being anticipated unpatentable over Frey et al. (U.S. Patent Number 6,029,168, hereinafter ‘Frey’) in view of Mukherjee (U.S. Patent Number 6,466,978, hereinafter ‘Mukherjee’). Appellant traverses this rejection for at least the following reasons. Different groups of claims are addressed under their respective subheadings.

Claim 1:

Claim 1 recites a method comprising, in part, associating a **signature** with the data object, **indicative of a state of the data object**; querying the second storage environment **for a change to the signature** in preparation for a data access operation on the data object, and **updating** the first data structure **if the signature has changed**.

In rejecting claim 1 in the Final Action, the Examiner asserts that Frey discloses “associating a signature with the data object”. The Examiner’s assertion is incorrect. The Examiner first acknowledges (on page 2 of the Final Action) that col. 2, lines 55 – 61, which disclose “unique pointers”, do not appear to address the limitation “indicative of the state of the data object”, but then proceeds to repeat the citation of col. 2, lines 55 – 61 in the rejection (see page 6 of the Final Action) and adds the citations of lines 17 – 31 of col. 8 to “provide the citation for the argued limitation”. The Examiner then cites Frey, col. 8, lines 4 – 17 as teaching “querying the second storage environment for a change to the signature”. Appellant respectfully disagrees with this assertion as well. Col. 2, lines 55 – 61 and col. 8, lines 4 – 31 of Frey are reproduced in full below:

It is another object of this invention to provide system-wide unique pointers for consistent, reliable access to distributed file data blocks. The present invention utilizes a specialized data element which enables indication of the location of another data element in a very diverse and distributed file data environment residing across several computing systems and storage devices. (Frey, col. 2, lines 55 - 61)

In the case where a file access manager 44, on a non-starter node, receives a request to retrieve blocks of a file block which does not exist, the file access manager 44 sends a query to the file access manager of the starter node for the file which the request is made. The query requests that the file access manager of the starter node resolve the conflict as to whether the file has just been created and the user is too early or the file has just been deleted and the user is too late. The file access manager of the starter node can respond to the query by (1) confirming that the file request is valid and so the requesting file access manager should allocate or write space for the file block, or (2) stating that the file block request is invalid as the file never existed or was deleted and so the file access manager should deny the request with an error showing that no such file exists. Thus, each file access manager only keeps file meta-data associate with files and file blocks that are stored on storage that it manages. This localization of file resource information contrasts with the global file manager methods shown in the prior art examples of FIG. 3.

If desired, parity blocks may be computed based on subsets of the file data blocks. It is to be emphasized that the parity block and its generation is optional and is based on a file parameter and is not limited by the position where any particular file data blocks are to be stored. Additional file parameters may include extended attribute file

parameters such as, for example, security access levels or encryption information.
(Frey, col. 8, lines 4 – 31)

In contrast to the Examiner's assertion, the additional lines 17 – 31 cited from col. 8 of Frey as shown above also do not teach or suggest a signature indicative of a state of the data object, where a second storage environment is queried for a change to the signature in preparation for a data access operation, as recited in claim 1. In Frey, the "query" sent to the file access manager of the starter node is to "resolve the conflict as to whether the file has just been created and the user is too early or the file has just been deleted and the user is too late", not for a "change to a signature". In the Final Action, the Examiner appeared to once again indicate (as first indicated in the Non-Final Office Action of May 19, 2005), e.g., by underlining the phrase "unique pointers" in the lines from col. 2, that it was the "unique pointers" of Frey that correspond to the signature of claim 1. However, the newly-recited lines from col. 8 (lines 17 – 31) of Frey are silent with respect to the "unique pointers", so if the Examiner was asserting that the "unique pointers" teach the "signature" of claim 1, the recitation of lines 17 – 31 of col. 8 of Frey does not appear to overcome the incorrectness of the original rejection.

In Appellant's response to the Final Action dated October 24, 2005, Appellant respectfully requested the Examiner, if the Examiner's intent in the Final Action was to withdraw the rejection on the basis of the "unique pointers" language, to indicate exactly where in Frey signatures with the features recited in claim 1 are taught (e.g., exactly which of the objects described in Frey have each of the features of the signature recited in claim 1). Appellant notes that in the Advisory Action mailed on November 21, 2005, the Examiner once again cites the same lines of Frey (col. 8, lines 4 – 17 and 17 – 31) in support of the rejection. However, none of the objects disclosed in the cited lines of col. 8, including "file meta-data", "parity blocks", and "file parameters", possess the combination of features of the signature recited in claim 1. Appellant can find no teaching or suggestion anywhere in either Frey or Mukherjee, taken singly or in combination, of such a signature.

Further with respect to claim 1, the Examiner asserts in the Final Action (and repeats in the Advisory Action) that Mukherjee discloses “updating the first data structure”, and cites Column 10, lines 31 – 38 and lines 57 – 59, and Column 5, lines 35 – 41 of Mukherjee in support. Once again, the Examiner does not appear to address the limitation of “updating the first data structure **if the signature has changed**”, as recited in claim 1. The lines of Mukherjee cited by the Examiner are:

(Col. 10, lines 31 – 38): Referring to FIG. 6C, a client 131 initiates a file operation by locating the file and then sending a request to the associated disk manager 134. At step 156 the client 131 queries the file locator table 144 to determine upon which disk 132 the desired file is located. The client 131 then queries the disk locator table 142 at step 158 to find the disk manager 134 that is associated with the disk 132, step 158, sends a file request to that disk manager 134, step 160.

(Col. 10, lines 51 – 59): The network manager 136 evaluates the request, step 168, and grants the request from the disk manager 134 if there is sufficient network bandwidth, step 170. After the request from the disk manager is granted, the network manager 136 and disk manager 134 update their corresponding bandwidth records to reflect the change to available bandwidth, step 172. The control information in the network manager mirror 140 and disk manager mirror 138 is then updated, step 174. The control information in the network manager mirror 140 and disk manager mirror 138 is then updated, step 174.

(Col. 5, lines 35 – 41): Metadata management is a medium load server task that is required in a serverless file system approach if the clients cache metadata information. When write operations by other clients cause the metadata to be modified, the client that is acting as server either logically remaps the cached metadata or informs the other clients to reload the new metadata.

The “queries” to the file locator table and the disk locator table in lines 31 – 38 are not about changes to a signature. The updating of the “control information” in lines 57 – 59 has nothing to do with a change to a signature, or with a query to determine whether a signature has changed. The metadata modification, remapping, and reloading of col. 5, lines 35 – 41 of Mukherjee is not related to a signature or to a query about a signature change. There is no teaching or suggestion in the cited lines of Mukherjee, or anywhere else in Frey or Mukherjee, of a query to determine whether a signature with the features recited in claim 1 has changed, and of updating a data structure **if the signature has changed**.

Still further with respect to claim 1, Appellant furthermore respectfully disagrees with the Examiner's assertion that it would have been obvious to modify the teachings of Frey with the teachings of Mukherjee "in order to maintain system consistency and accuracy especially when subsequent requests are received." Mukerjee teaches a distributed file server system in which "to provide for recovery from the failure of a server or client, mirrors of the disk managers associated with each of the configurations are maintained" (see, e.g., Mukherjee, col. 9, lines 19 – 24). Frey specifically teaches that "the present invention provides a distribution of entries, such as fields and indexes or any commands that are used to describe and work with data, without replication of the entries" (Col. 2, lines 47 – 50). Accordingly, Frey actually appears to teach away from a system that employs disk manager mirrors as disclosed by Mukherjee. Furthermore, Frey already provides "system-wide unique pointers for consistent, reliable access to distributed file data blocks" (Frey, col. 2, lines 55 – 58), so Frey's system does not appear to need any modification "in order to maintain system consistency and accuracy" as alleged by the Examiner. Appellant respectfully submits that the Examiner has not provided evidence of sufficient motivation for combining the teachings of Mukherjee with those of Frey.

Thus, for at least the reasons presented above, the rejection of claim 1 is not supported by the cited art and removal thereof is respectfully requested.

Claim 2:

Claim 2 depends on claim 1. Since the rejection of claim 1 has been shown to be in error, claim 2 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 3:

With respect to claim 3, which depends on claim 1, the Examiner asserts that Frey teaches the additional limitation "an operating system of the first storage environment

does not support the second storage environment”, citing col. 2, line 57 – col. 3, line 3 and col. 3, lines 43 – 53 of Frey. The Examiner is incorrect. While Frey teaches, “a very diverse and distributed file data environment residing across several computing systems and storage devices” (col. 2, lines 55 – 62), Frey does **not** teach, in the cited lines or anywhere else, that an operating system of a first storage environment does not support a second storage environment, as recited in claim 3, where the first and second storage environments have the features recited in parent claim 1.

For at least the reasons presented above, the rejection of claim 3 is not supported by the cited art and removal thereof is respectfully requested.

Claim 4:

Claim 4 depends on claim 1. Since the rejection of claim 1 has been shown to be in error, claim 4 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 5:

With respect to claim 5, the Examiner asserts that Mukherjee discloses “wherein the generation further includes detecting a mirroring of the data object on at least two storage devices within the second storage environment”, and cites col. 9, lines 19 – 45 and col. 16, lines 21 – 34 of Mukherjee in support. The Examiner is incorrect. The cited portions of Mukherjee teach “mirrors of the managers” (e.g., “network managers, file managers and cluster managers”) and describe the functionality of these mirrored managers, but are silent with respect to “**detecting** a mirroring of a data object” as part of a “**generation of a data structure** from the reference representing one or more physical locations of the data object within the second storage environment, wherein one or more extents of the data object within the second storage environment are provided during the generation”, as recited in the parent claims of claim 5.

For at least the reasons presented above, the rejection of claim 5 is not supported by the cited art and removal thereof is respectfully requested.

Claim 6:

Claim 6 depends on claim 1. Since the rejection of claim 1 has been shown to be in error, claim 6 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 7:

With respect to claim 7, the Examiner asserts that Frey discloses “wherein in retaining the first data structure the first data structure is validated with one or more subsequent references made to access the first data object”, citing col. 8, lines 4 – 17 of Frey in support. The Examiner is incorrect. The cited lines of Frey are reproduced above in the discussion regarding claim 1. In the cited lines, Frey teaches a “query” that “requests that the file access manager of the starter node resolve the conflict as to whether the file has just been created and the user is too early or the file has just been deleted and the user is too late.” The file access manager of the starter node can respond to the query by “(1) confirming that the file request is valid and so the requesting file access manager should allocate or write space for the file block, or (2) stating that the file block request is invalid as the file never existed or was deleted and so the file access manager should deny the request with an error showing that no such file exists.” However, contrary to the Examiner’s assertion (see page 9 of the Final Action, underlined section of cited lines of Frey, “to whether (validate)”), “resolving the conflict as to whether the file has just been created and the user is too early or the file has just been deleted and the user is too late” is very different from “validating the first data structure”, where the first data structure has the features recited in parent claim 1. In addition, determining whether or not “the file request” is valid is also different from validating a data structure in a first storage environment that represents one or more

physical locations of the data object within the second storage environment, as recited in claim 7 together with parent claim 1.

For at least the reasons presented above, the rejection of claim 7 is not supported by the cited art and removal thereof is respectfully requested.

Claim 8:

Claim 8 depends on claim 1. Since the rejection of claim 1 has been shown to be in error, claim 8 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 9:

Claim 9 recites, in part, **assembling a hierarchical map** representing a path to one or more storage locations of a data storage object, associating a **signature** with the map **indicative of a state of the data storage object, querying the signature for changes** in preparation for a data access operation on the data storage object, and **updating the map if the signature has changed**.

In the Final Action, the Examiner rejected claim 9 citing the same sections of Frey and Mukherjee that were relied upon in the rejection of claim 1. As discussed above in detail with respect to claim 1, neither Frey nor Mukherjee, taken singly or in combination, teach “a signature indicative of a state of a data storage object” that is queried for changes in preparation for a data access operation, much less associating such a signature with a map representing a path to one or more storage locations of the data storage object, or updating the map if the signature has changed, as recited in claim 9.

Further with respect to claim 9, the Examiner does not directly address the limitation of “**assembling a hierarchical map representing a path** to one or more storage locations of a data storage object” in the Final Action. The Examiner cites col. 9,

lines 15 – 22 and col. 5, lines 4 – 10 of Frey as teaching “generating a first data structure from the reference representing one or more physical locations”, but this limitation of claim 1 is clearly not the same as the limitation “assembling a hierarchical map representing a path to one or more storage locations” of claim 9. In col. 5, lines 4 – 10 Frey teaches that “starter information describing how a file is to spread over a node” is stored “in a table at a unique starter information index” and that “it is often efficient to include the mapping of file blocks to be stored on this node to their disk address with the starter information”. However, the cited lines of Frey are silent with respect to assembling a “hierarchical map” as recited in the claim. A “hierarchical map” may be represented using a variety of data structures, but this does not mean that any data structure (such as the table of Frey col. 5) “describing how a file is to spread over a node” or “including a mapping of file blocks to disk addresses” represents a hierarchical map. Neither Frey nor Mukherjee, taken singly or in combination, teach or suggest assembling a hierarchical map, where the map is associated with a signature, with the combination of features recited in claim 9.

Appellant also respectfully submits that the arguments regarding insufficient evidence of a motivation to combine the teachings of Mukherjee with those of Frey, provided above with respect to claim 1, apply with equal force to claim 9.

For at least the reasons presented above, the rejection of claim 9 is not supported by the cited art and removal thereof is respectfully requested.

Claim 10:

Claim 10 depends on claim 9. Since the rejection of claim 9 has been shown to be in error, claim 10 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 11:

Claim 11 depends on claim 10. Since claim 10 is believed to be in condition for allowance, claim 11 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 12:

Claim 12 depends on claim 9. Since the rejection of claim 9 has been shown to be in error, claim 12 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 13:

Claim 13, which depends on claim 9, recites the limitation “wherein the method is repetitively processed for one or more additional data objects residing in the first storage environment”. The Examiner does not appear to address the specific limitation of claim 13 in the Final Action. Instead, the Examiner cites lines col. 8, lines 24 – 31 of Frey (reproduced above in conjunction with Appellant’s argument with respect to claim 1) in rejecting claim 13 together with claim 4. The cited lines describe “parity blocks” and “file parameters”, but have nothing to do with the limitations of claim 13.

For at least the reasons presented above, the rejection of claim 13 is not supported by the cited art and removal thereof is respectfully requested.

Claim 14:

In regard to claim 14, the Examiner asserts that Frey discloses “wherein the method is used to create an image or copy of the first storage environment in the second storage environment”, and refers to “Figure 3B, Element No. 22, i.e., objects {A, B and E} in server 36 and the copy or image of the same in server 38” in support of this assertion. Once again, the Examiner is mistaken. Figure 3B does not show a “copy or image” of objects A, B or E. Rather, Figure 3B shows “files A, B, C, D and E striped

across two server nodes, a first server node 36 and a second server node 38” (col. 4, lines 39 – 41 of Frey). The letters A, B, E, etc. at the different server nodes 36 and 38 refer to different blocks of the files, not to copies or images of the files: see, e.g., col. 4, lines 24 – 26: “Each data storage system 22 (here, a disk) is shown divided into **stored file blocks labeled according to the file A, B, C, D or E from which the blocks originated**”. Furthermore, there is no teaching or suggestion in either Frey or Mukherjee of **using the method** recited in claims 9 and 13 (on which claim 14 depends) to create an image or copy of the first storage environment in the second storage environment.

For at least the reasons presented above, the rejection of claim 14 is not supported by the cited art and removal thereof is respectfully requested.

Claim 15:

Claim 15 recites a computer readable medium comprising, in part, “a signature representing a data object residing on a second computer readable medium, wherein the **signature is indicative of a state of the data object**”. The computer readable medium also comprises a map including a plurality of nodes: a first node representing the data object, a file system node representing a file system on a second computer readable medium, a volume node representing a volume manager associated with the file system, one or more partition nodes managed by the volume manager, and one or more disk identifications representing one or more storage devices housing the data object. The map is updated when a change to the signature is detected.

In the Final Action, the Examiner rejected claim 15 on the basis of the same portions of Frey and Mukherjee that were relied upon to reject claims 1 and 9. As pointed out earlier with respect to claim 1, neither Frey nor Mukherjee teach “a signature indicative of a state of a data object”. Therefore, neither Frey nor Mukherjee can teach that a computer readable medium that comprises such a signature also includes a map with the specific features recited in claim 15. In fact, the Examiner does not even point out where in the prior art a map comprising the combination of specific nodes (e.g., a first

node representing the data object, a file system node, a volume node, one or more partition nodes, etc.) recited in claim 15 is taught or suggested.

Appellant also respectfully submits that the arguments regarding insufficient evidence of a motivation to combine the teachings of Mukherjee with those of Frey, provided above with respect to claim 1, apply with equal force to claim 15.

For at least the reasons presented above, the rejection of claim 15 is not supported by the cited art and removal thereof is respectfully requested.

Claim 16:

With respect to claim 16, the Examiner asserts that Mukherjee teaches the limitation “wherein the map is represented as a tree data structure on the computer readable media” in col. 15, lines 5 - 13. The Examiner is again mistaken. In the cited lines, Mukherjee teaches that “during the cluster expansion process the initial load of a newly created file manager is determined in a binary tree fashion. The list of files in the cluster 401 is statically divided and allocated such that all the leaves, file managers 406, of the tree have a fixed set of files to manage. The file list and tree configuration is provided to a client 408 when it enters the hybrid file system 400. A client 408 that knows the tree configuration can derive which files are managed by a selected file manager 404.” However, a “tree configuration” indicating “which files are managed by a selected file manager” is very different from representing a map with the specific features recited in parent claim 15 as a tree, where the map includes respective nodes for a file system, a volume manager, and one or more partitions.

For at least the reasons presented above, the rejection of claim 16 is not supported by the cited art and removal thereof is respectfully requested.

Claim 17:

Claim 17 (in combination with its parent claim 15) recites a map comprising, in pertinent part, “**a first node representing the data object, a file system node representing a file system on a second computer readable medium, a volume node representing a volume manager associated with the file system, and one or more partition nodes managed by the volume manager, wherein each node includes metadata**”. The Examiner asserts in the Final Action that the additional limitation that “each node includes metadata” is taught by Mukherjee (col. 5, lines 35 – 41, col. 7, lines 6 – 7) and also by Frey (Abstract, Figures 7 – 9 and corresponding text). The Examiner is incorrect. Claim 17 is directed to **metadata included in each node of a map**, where the map comprises **a node representing a data object, a node representing a file system, a node representing a volume manager, and one or more nodes representing partitions**. Appellant can find no teaching or suggestion of such a map, **where each node of the map includes metadata**, anywhere in Mukherjee or Frey.

For at least the reasons presented above, the rejection of claim 17 is not supported by the cited art and removal thereof is respectfully requested.

Claim 18:

With respect to claim 18, the Examiner asserts that the limitation “wherein the map is used by an accessing set of executable instructions having access to a second file system which is incompatible with the first file system” is taught in col. 2, line 57 – col. 3, line 3 and col. 3, lines 43 – 53 of Frey. The Examiner is incorrect. While Frey teaches “a very diverse and distributed file data environment residing across several computing systems and storage devices” (col. 2, lines 55 – 62), Frey does **not** teach, in the cited lines or anywhere else, that a map with the features recited in parent claim 15 “is used by an accessing set of executable instructions having access to a second file system which is incompatible with the first file system”.

For at least the reasons presented above, the rejection of claim 18 is not supported by the cited art and removal thereof is respectfully requested.

Claim 19:

Claim 19 depends on claim 18. Since claim 18 is believed to be in condition for allowance, claim 19 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 20:

Claim 20 depends on claim 19. Since claim 19 is believed to be in condition for allowance, claim 20 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 21:

Claim 21 recites, in part, a first and a second file system, a map generated to represent one or more physical locations for each of one or more data objects of the second file system and used by the first file system when one or more of the data objects are referenced, and a **signature indicative of the state** of the one or more data objects, where the **map is updated when changes are detected and associated with the signature**.

In the Final Action, the Examiner relies on the same prior art to reject claim 21 that was relied upon in the rejection of claims 1, 9 and 15. The Examiner fails to point out specifically where in Frey or Mukherjee **a first and a second file system** with the features recited in claim 21 are taught. Furthermore, as pointed out earlier with respect to claim 1, neither Frey nor Mukherjee, taken singly or in combination, teach **a signature indicative of a state of one or more data objects**. Therefore, the combination of Frey and Mukherjee clearly cannot teach that a map representing physical locations of data objects in a file system is updated when changes are detected and associated with such a signature, as recited in claim 21.

Appellant also respectfully submits that the arguments regarding insufficient evidence of a motivation to combine the teachings of Mukherjee with those of Frey, provided above with respect to claim 1, apply with equal force to claim 21.

Thus, for at least the reasons presented above, the rejection of claim 21 is not supported by the cited art and removal thereof is respectfully requested.

Claim 22:

Claim 22 depends on claim 21. Since the rejection of claim 21 has been shown to be in error, claim 22 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 23:

Claim 23 depends on claim 21. Since the rejection of claim 21 has been shown to be in error, claim 23 is also believed to be in condition for allowance, and the removal of the rejection thereof is respectfully requested.

Claim 24:

With respect to claim 24, which depends on claims 23 and 21, the Examiner asserts that Frey teaches the additional limitation “wherein the file systems operate within different operating systems”, citing col. 2, line 57 – col. 3, line 3 and col. 3, lines 43 – 53 of Frey. The Examiner is incorrect. While Frey teaches “a very diverse and distributed file data environment residing across several computing systems and storage devices” (col. 2, lines 55 – 62), Frey does **not** teach, in the cited lines or anywhere else, that a first and a second file system operate within different operating systems, as recited in claim 24, where the first and second file systems have the features recited in parent claim 21.

For at least the reasons presented above, the rejection of claim 24 is not supported by the cited art and removal thereof is respectfully requested.

Claim 25:

Claim 25 recites a map that is used to “**replicate the second file system within the first file system in a first file system format**”. In rejecting claim 25, in the Final Action the Examiner asserts that Mukherjee “discloses different data or file formats” and cites column 1, lines 27 – 29 and column 7, lines 63 – 67 of Mukherjee in support. The Examiner is incorrect. The cited lines of Mukherjee are:

(Column 1, lines 27 – 29): Increased constraints are placed on the storage and retrieval of multimedia data over traditional textual and numeric data due to inherent differences in the characteristics of the data types.

(Column 7, lines 63 – 67): At step 90 an initial quantity of bandwidth is allocated to each manager. The initial bandwidth allocations are based on criteria such as prior experience, the expected workload, and the type of data the manager controls.

The cited lines do **not** teach or suggest “**using a map to replicate a second file system within a first file system in a first file system format**”, as recited in claim 25. Neither Frey nor Mukherjee, taken singly or in combination, teach the limitations recited in claim 25.

For at least the reasons presented above, the rejection of claim 25 is not supported by the cited art and removal thereof is respectfully requested.

Second Ground of Rejection:

Claim 26 stands finally rejected under 35 U.S.C. 103(a) as being unpatentable over Frey in view of Mukherjee, further in view of Banerjee et al., (U.S. Patent No. 6, 795, 830, hereinafter, ‘Banerjee’).

With respect to claim 26, the Examiner asserts that “Banerjee discloses XML data structure and distributing the data structure through the Internet” and refers to “col. 23, lines 1 – 31, i.e., the components defined in the XML associated with the template are added to the customer site file”. (As Appellant noted in response to an earlier Non-Final Office Action, Appellant assumes the Examiner meant to cite column 22, not column 23, which does not contain the text specifically cited by the Examiner.) Banerjee teaches a wizard used to build web sites using stored web site components: “The web site building wizard appliance stores a large number of such components so that a novice user does not have to reinvent them. The components can be represented in any manner known in the art. In one embodiment, each component is described by an extensible markup language (XML) document.” (Column 21, lines 29 – 35). The “components” taught by Banerjee “handle specific functional needs of an enterprise operating a web site on the Internet.” (Column 20, lines 65 – 66). Examples of the components include “a site logo, a site name, a legal notice, terms of use, and a privacy statement” and “a product or service description, a product/service price, a list of products/services, a component for searching for a product/service among the list of products, a hierarchical list of products/services within categories and subcategories, a component for searching products/services falling within a category or subcategory, support contact information, personnel lists, an item to search a personnel list, a map of the web site, links to related web sites, a calendar of appointments/events, a banner advertisement and a shopping cart.” (Column 21, lines 3 – 20).

In rejecting claim 26, the Examiner specifically asserts that Banerjee’s teaching that the “template are added to the customer site XML file” is a disclosure of “distributing the data structure through the Internet”. The Examiner is incorrect in this assertion. Banerjee **does not teach or suggest distributing anything through the Internet in adding to the customer site XML file**: see, e.g., column 22, lines 22 – 32: “According to one embodiment, the wizard represents the customer's web site using an XML file (the "customer site XML file"). When the customer initially selects a template, the components defined in the XML associated with the template are added to the customer site XML file. As the user goes from screen to screen in the wizard, the user

specifies changes, deletions or additions to the components in the site. In response to the user input entered through these screens, the wizard changes, deletes or adds XML content in the customer site XML file.”

Further, the Examiner asserts that “it would have been obvious to modify both Frey and Mukherjee in view of Banerjee **due to the wide use of the XML language** especially when the Internet is being used as the protocol for the transfer; and also to **increase system performance** by using one very popular language that would minimize the formatting and re-formatting of data to only one common format”. Appellant respectfully disagrees. An alleged “wide use of the XML language” does not make using XML to generate a portable representation of a data structure **representing one or more physical locations of a data object within a second storage environment, and distributing the portable representation to a third storage environment**, as recited in claim 26, obvious. Neither Banerjee, Mukherjee nor Frey, taken singly or in combination, teach or suggest creating or distributing XML representations of such a data structure. Furthermore, there is no teaching or suggestion of the performance impact of formatting or re-formatting a representation of such a data structure anywhere in Banerjee, Mukherjee or Frey.

For at least the reasons presented above, the rejection of claim 26 is not supported by the cited art and removal thereof is respectfully requested.

VIII. CONCLUSION

For the foregoing reasons, it is submitted that the Examiner’s rejection of claims 1-37 was erroneous, and reversal of his decision is respectfully requested.

The Commissioner is authorized to charge the appeal brief fee of \$500.00 and any other fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5760-16700/BNK. This Appeal Brief is submitted with a return receipt postcard.

Respectfully submitted,



B. Noël Kivlin
Reg. No. 33,929
Attorney for Appellants

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
(512) 853-8850

Date: February 21, 2006

IX. CLAIMS APPENDIX

The claims on appeal are as follows.

1. A method to map a storage environment data object, comprising;
receiving a reference to the data object in a first storage environment, wherein the
data object resides in a second storage environment;
generating a first data structure from the reference representing one or more
physical locations of the data object within the second storage
environment;
associating a signature with the data object, indicative of a state of the data object;
retaining the first data structure in the first storage environment;
querying the second storage environment for a change to the signature in
preparation for a data access operation on the data object;
updating the first data structure if the signature has changed; and
performing the data access operation using the first data structure to interface with
one or more of the physical locations of the data object from the first
storage environment.
2. The method of claim 1, wherein in retaining the first data structure one or more
additional references access the data object using the first data structure.
3. The method of claim 1, wherein receiving the reference an operating system of
the first storage environment does not support the second storage environment.
4. The method of claim 1, wherein during generation one or more extents of the data
object within the second storage environment are provided.
5. The method of claim 4, wherein the generation further includes detecting a
mirroring of the data object on at least two storage devices within the second storage
environment.

6. The method of claim 1, wherein during generation metadata associated with the second storage environment and the data object are provided.
7. The method of claim 1, wherein in retaining the first data structure the first data structure is validated with one or more subsequent references made to access the data object.
8. The method of claim 1, wherein the method is used to interface a first database using the first storage environment with a second database using the second storage environment.
9. A method to represent a data storage object, comprising:
 - identifying one or more storage locations for the data storage object housed within a first storage environment;
 - assembling a hierarchical map representing a path to one or more of the storage locations;
 - associating a signature with the map indicative of a state of the data storage object;
 - querying the signature for changes in preparation for a data access operation on the data storage object;
 - updating the map if the signature has changed; and
 - using the map in a second storage environment to access the data storage object.
10. The method of claim 9, wherein while assembling the map attribute data are acquired and associated with the first storage environment.
11. The method of claim 10, wherein assembling further includes acquiring attribute data associated with the data storage object.

12. The method of claim 9, wherein in identifying one or more of the storage locations, the data storage object is identified as at least one of a file system, a file, a database, a volume, and a portion of data within a file.
13. The method of claim 9, wherein the method is repetitively processed for one or more additional data objects residing in the first storage environment.
14. The method of claim 13, wherein the method is used to create an image or copy of the first storage environment in the second storage environment.
15. A first computer readable medium having a data map and a signature representing a data object residing on a second computer readable medium, wherein the signature is indicative of a state of the data object, the map comprising:
- a first node representing the data object;
 - a file system node representing a file system on the second computer readable medium;
 - a volume node representing a volume manager associated with the file system;
 - one or more partition nodes managed by the volume manager;
 - one or more disk identifications representing one or more storage devices housing the data object; and
- wherein the map is updated when a change to the signature is detected.
16. The map of claim 15, wherein the map is represented as a tree data structure on the computer readable media.
17. The map of claim 15, wherein each node includes metadata.
18. The map of claim 15, wherein the map is used by an accessing set of executable instructions having access to a second file system which is incompatible with the first file system.

19. The map of claim 18, wherein the data object is referenced and modified by the accessing set of executable instructions from the second file system.

20. The map of claim 19, wherein the map is updated if one or more values associated with one or more of the nodes or identifications are modified.

21. A storage environment system, comprising:

- a first file system;

- a second file system;

- one or more data objects residing in the second file system and capable of being referenced within the first file system;

- a map generated to represent one or more physical locations for each of the one or more data objects and used by the first file system when one or more of the data objects are referenced; and

- a signature indicative of a state of the one or more data objects;

- wherein the map is updated when changes are detected and associated with the signature;

22. The system of claim 21, wherein the map is modified when one or more of the physical locations for each of the one or more data objects changes.

23. The system of claim 21, wherein the file systems reside in different computing environments.

24. The system of claim 23, wherein the file systems operate within different operating systems.

25. The system of claim 24, wherein the map is used to replicate the second file system within the first file system in a first file system format.

26. The method of claim 1, further comprising:

creating a portable representation of the data structure using extensible markup language (XML); and
distributing the portable representation to a third storage environment over the Internet.

X. EVIDENCE APPENDIX

No evidence submitted under 37 CFR §§ 1.130, 1.131 or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

XI. RELATED PROCEEDINGS APPENDIX

There are no related proceedings.

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☒ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.